

What is claimed is:

- 1 1. A method of communicating between nodes in a clustered computer
2 system, the method comprising:
 - 3 (a) communicating a port identifier from a first node to a second node
4 coupled to the first node over a point-to-point network, wherein the first node
5 includes a plurality of network ports and a plurality of communication
6 registers, wherein each communication register is dedicated to an associated
7 network port among the plurality of network ports and is configured to store
8 data received over such associated network port, and wherein the port
9 identifier identifies a network port among the plurality of network ports to
10 which the second node is coupled to the first node; and
 - 11 (b) communicating data from the second node to the first node by
12 initiating a write operation on the first node using the second node to store the
13 data in the communication register associated with the network port identified
14 by the port identifier.
- 1 2. The method of claim 1, further comprising detecting in the first node the
2 storage of data in the communication register associated with the network port
3 identified by the port identifier.
- 1 3. The method of claim 2, further comprising generating an interrupt on the
2 first node in response to detecting the storage of data in the communication register
3 associated with the network port identified by the port identifier.
- 1 4. The method of claim 3, further comprising processing the interrupt by
2 processing the data stored in the communication register associated with the network
3 port identified by the port identifier, and clearing the interrupt.
- 1 5. The method of claim 4, wherein detecting the storage of data comprises
2 detecting a non-zero value stored in any of the plurality of communication registers,

3 and wherein clearing the interrupt comprises resetting the plurality of communication
4 registers to zero values.

1 6. The method of claim 1, wherein communicating the data comprises
2 sequentially storing a plurality of commands in the communication register associated
3 with the network port identified by the port identifier, the method further comprising
4 processing each of the plurality of commands in the first node.

1 7. The method of claim 1, further comprising initiating, with the second node,
2 a read operation for a configuration register in the first node, wherein communicating
3 the node identifier is performed in response to the read operation.

1 8. The method of claim 1, wherein communicating the node identifier is
2 performed in response to a read request sent over the point-to-point network by the
3 second node.

1 9. The method of claim 1, wherein the plurality of communication registers
2 are allocated a range of register addresses in a register address space for the node, and
3 wherein communicating the data comprises sending a write request to the register
4 address of the communication register associated with the network port identified by
5 the port identifier.

1 10. A circuit arrangement, comprising:

2 (a) a plurality of network ports, each configured to couple a first node
3 from a clustered computer system to another node in the clustered computer
4 system over a point-to-point network;

5 (b) a plurality of communication registers, each dedicated to an
6 associated network port among the plurality of network ports and configured
7 to store data received through such associated network port; and

8 (c) a control circuit coupled to the plurality of communication
9 registers and configured to automatically notify the first node in response to
10 storage of data in any of the plurality of communication registers.

1 11. The circuit arrangement of claim 10, wherein the control circuit is
2 configured to detect the storage of data in a communication register among the
3 plurality of communication registers by detecting a non-zero value stored in such
4 communication register.

1 12. The circuit arrangement of claim 11, wherein the control circuit is
2 configured to automatically notify the first node by generating an interrupt.

1 13. The circuit arrangement of claim 12, wherein the control circuit is
2 configured to generate a common interrupt for all of the plurality of communication
3 registers.

1 14. The circuit arrangement of claim 10, wherein each communication register
2 includes a plurality of binary outputs, and wherein the control circuit comprises at
3 least one logic gate configured to generate an interrupt signal by performing a logical-
4 OR operation on all of the binary outputs of the plurality of communication registers.

1 15. The circuit arrangement of claim 10, wherein the control circuit is further
2 configured to output a port identifier over a first network port among the plurality of
3 network ports in response to a read request received over the first network port, the

4 port identifier identifying the first network port as the network port from which the
5 read request was received.

1 16. The circuit arrangement of claim 15, further comprising a configuration
2 register, wherein the control circuit is configured to output data stored in the
3 configuration register in response to the read request.

1 17. The circuit arrangement of claim 16, wherein the plurality of
2 communication registers are allocated a range of register addresses in a register
3 address space for the node, and wherein the control circuit is configured to store data
4 received over a first network port among the plurality of network ports in the
5 communication register associated with the first network port in response to a write
6 request addressed to the register address of the communication register associated
7 with the network port identified by the port identifier.

1 18. An integrated circuit device comprising the circuit arrangement of claim
2 10.

1 19. A program product comprising a hardware definition program that defines
2 the circuit arrangement of claim 10, and a signal bearing medium bearing the
3 hardware definition program, wherein the signal bearing medium includes at least one
4 of a transmission medium and a recordable medium.

- 1 20. A node for use in a clustered computer system, the node comprising:
- 2 (a) a plurality of network ports, each configured to couple to another
- 3 node in the clustered computer system over a point-to-point network;
- 4 (b) a plurality of communication registers, each dedicated to an
- 5 associated network port among the plurality of network ports and configured
- 6 to store data received through such associated network port; and
- 7 (c) a control circuit coupled to the plurality of communication
- 8 registers and configured to automatically notify the node in response to storage
- 9 of data in any of the plurality of communication registers.
- 1 21. The node of claim 20, wherein the control circuit is configured to generate
- 2 the notification by signaling an interrupt in response to any of the plurality of
- 3 communication registers storing a non-zero value.

PROVISIONAL PATENT APPLICATION

- 1 22. A clustered computer system comprising:
- 2 (a) a plurality of nodes, each node including:
- 3 (i) a plurality of network ports;
- 4 (ii) a plurality of communication registers, each dedicated to an
5 associated network port among the plurality of network ports and
6 configured to store data received through such associated network port;
7 and
- 8 (iii) a control circuit coupled to the plurality of communication
9 registers and configured to automatically notify such node in response
10 to storage of data in any of the plurality of communication registers;
11 and
- 12 (b) a plurality of point-to-point network interconnects, each coupled
13 between a pair of nodes from the plurality of nodes through network ports on
14 each of the pair of nodes.

IBM ROCKWELL COMMUNICATIONS